arXiv:1401.4744v2 [cs.SC] 2 May 2014

Sparse interpolation over finite fields via low-order roots of unity

Andrew Arnold
Cheriton School of Computer Science
University of Waterloo
a4arnold@uwaterloo.ca

Mark Giesbrecht
Cheriton School of Computer Science
University of Waterloo
mwg@uwaterloo.ca

Daniel S. Roche
Computer Science Department
United States Naval Academy
roche@usna.edu

May 5, 2014

Abstract

We present a new Monte Carlo algorithm for the interpolation of a straight-line program as a sparse polynomial f over an arbitrary finite field of size q. We assume *a priori* bounds D and T are given on the degree and number of terms of f. The approach presented in this paper is a hybrid of the diversified and recursive interpolation algorithms, the two previous fastest known probabilistic methods for this problem. By making effective use of the information contained in the coefficients themselves, this new algorithm improves on the bit complexity of previous methods by a "soft-Oh" factor of T, $\log D$, or $\log q$.

1 Introduction

Let $\mathbb{F}_q$ be a finite field of size q and consider a "sparse" polynomial

$$f = \sum_{i=1}^{t} c_i z^{e_i} \in \mathbb{F}_q[z], \tag{1.1}$$

where the $c_1, \ldots, c_t \in \mathbb{F}_q$ are nonzero and the exponents $e_1, \ldots, e_t \in \mathbb{Z}_{\geq 0}$ are distinct. Suppose f is provided as a straight-line program, a simple branch-free program which evaluates f at any point (see below for a formal definition). Suppose also that we are given an upper bound D on the degree of f and an upper bound T on the number of non-zero terms t of f. Our goal is to recover the standard form (1.1) for f, that is, to recover the coefficients c_i and their corresponding exponents e_i as in (1.1). Our main result is as follows.

Theorem 1.1. *Let $f \in \mathbb{F}_q[z]$ with at most T non-zero terms and degree at most D, and let $0 < \epsilon \leq 1/2$. Suppose we are given a division-free straight-line program $\mathcal{S}_f$ of length L that computes f. Then there exists an algorithm (presented below) that interpolates f, with probability at least $1 - \epsilon$, with a cost of*

$$\widetilde{\mathcal{O}}\left(LT \log^2 D \left(\log D + \log q\right) \log \tfrac{1}{\epsilon}\right)^{\dagger}$$

bit operations.

This cost improves on previous methods by a factor of T, $\log D$, or $\log q$, and may lay the groundwork for even further improvements. See Table 1 below for a detailed comparison of the complexity of various approaches to this problem.

1.1 Straight-line programs and sparse polynomials

The interpolation algorithm presented in this paper is for straight-line programs, though it could be adapted to other more traditional models of interpolation. Informally, a straight-line program is a very simple program, with no branches or loops, which evaluates a polynomial at any point, possibly in an extension ring or field. Straight-line programs serve as a very useful model to capture features of the complexity of algebraic problems (see, e.g., (Strassen, 1990)) especially with respect to evaluation in extension rings, as well as having considerable efficacy in practice (see, e.g., (Freeman, Imirzian, Kaltofen, and Yagati, 1988)).

More formally, a division-free *Straight-Line Program* over a ring R, henceforth abbreviated as an *SLP*, is a branchless sequence of arithmetic instructions that represents a polynomial function. It takes as input a vector $(a_1, \ldots, a_K)$ and outputs a vector $(b_1, \ldots, b_L)$ by way of a series of instructions $\Gamma_i : 1 \leq i \leq L$ of the form $\Gamma_i : b_i \longleftarrow \alpha \star \beta$, where $\star$ is an operation $'+', '-'$, or $'\times'$, and $\alpha, \beta \in \mathsf{R} \cup \{a_1, \ldots, a_K\} \cup \{b_0, \ldots, b_{i-1}\}$. The inputs and outputs may belong to R or a ring extension of R. We say a straight-line program *computes* $f \in \mathsf{R}[x_1, \ldots, x_K]$ if it sets b_L to $f(a_1, \ldots, a_K)$.

The straight-line programs in this paper compute over finite fields $\mathbb{F}_q$ with q elements, and ring extensions of $\mathbb{F}_q$. We assume that elements of $\mathbb{F}_q$ are stored in some reasonable representation with $O(\log q)$ bits, and that each field operation requires $\widetilde{\mathcal{O}}(\log q)$ bit operations. Similarly, we assume that elements in a field extension $\mathbb{F}_{q^s}$ of $\mathbb{F}_q$ can be represented with $O(s \log q)$ bits, and that operations in $\mathbb{F}_{q^s}$ require $\widetilde{\mathcal{O}}(s \log q)$ bit operations.

Each of the algorithms described here determines f by *probing* its SLP: executing it at an input of our choosing and observing the output. This is analogous to evaluation of the polynomial at a point in an extension of the ground field in the more traditional interpolation model. To fairly account for the cost of such a probe we define the *probe degree* as the degree of the ring extension over $\mathbb{F}_q$ in which the probe lies. A probe of degree u costs $\widetilde{\mathcal{O}}(Lu)$ field

†For functions $\phi, \psi : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$, we say $\phi \in \widetilde{\mathcal{O}}(\psi)$ if and only if $\phi \in \mathcal{O}(\psi(\log \psi)^c)$ for a constant $c \geq 0$.

operations, or $\widetilde{\mathcal{O}}(Lu \log q)$ bit operations. The *total probe size* is the sum of the probe degrees of all the probes used in a computation.

Polynomials are also stored with respect to the power basis (powers of x) in a *sparse representation*. In particular, f as in (1.1) would be stored as a list $[(c_i, e_i)]_{1 \le i \le t}$. Given such a polynomial f, we let $\texttt{coeff}(f, k)$, denote the coefficient of the term of f of degree k (which may well be zero). We let $\texttt{exponents}(f)$ denote the sorted list of exponents of the nonzero terms in f, and $\#f$ denote the number of nonzero terms of f. When we write $r(z) = f(z) \bmod d(z)$, we assume that the modular image r is reduced, i.e., it is stored as the remainder of f divided by d, and $\deg r < \deg d$. We frequently refer to such r as "images" of the original polynomial f, as they reveal some limited amount of information about f.

The problem of (sparse) interpolation can then be seen as one of conversion between representations: Efficiently transform a polynomial given by a straight-line program into a "standard" sparse representation with respect to a power basis in x as in (1.1), given "size" constraints D and T as above.

1.2 Previous results

The following table gives a comparison of existing algorithms for the sparse interpolation of straight-line programs.

	Bit complexity	Type
Dense	$LD \log q$	Det
Garg & Schost	$LT^4 \log^2 D \log q$	Det
LV G & S	$LT^3 \log^2 D \log q$	LV
Diversified	$LT^2 \log^2 D(\log D + \log q)$	LV
Recursive	$LT \log^3 D \log q \cdot \log \frac{1}{\epsilon}$	MC
This paper	$LT \log^2 D(\log D + \log q) \cdot \log \frac{1}{\epsilon}$	MC

Table 1: *A comparison of interpolation algorithms for straight-line programs (ignoring polylogarithmic factors).*
Det=Deterministic, LV=Las Vegas, MC=Monte Carlo.

Most of these are probabilistic algorithms. By a *Las Vegas* algorithm we mean one that runs in the expected time stated but produces an output of guaranteed correctness. A *Monte Carlo* algorithm takes an additional parameter $\epsilon \in (0, 1)$ and produces an output guaranteed to be correct with probability at least $1 - \epsilon$.

The algorithm of Garg and Schost (2009) finds a *good prime*, that is, a prime that separates all the terms of f. For f given by (1.1),

$$f \bmod (z^p - 1) = \textstyle\sum_{i=1}^{t} c_i z^{e_i \bmod p}, \tag{1.2}$$

and so the terms of the *good image* $f \bmod (z^p - 1)$ remain distinct provided the exponents e_i are distinct modulo p. This good image gives us $t = \#f$,

which in turn makes it easy to identify other good primes. Their algorithm then constructs the symmetric polynomial $\chi(y) = \prod_{i=1}^{t}(y - e_i)$ by Chinese remaindering of images $\chi(y) \bmod p_j$, for sufficiently many good primes p_j, $1 \leq j \leq \ell$. Note that the image $f \bmod (z^{p_j} - 1)$ gives the values $e_i \bmod p_j$ and hence $\chi(y) \bmod p_j$. Given χ, the algorithm then factors χ to obtain the exponents e_i. The corresponding coefficients of f may be obtained by inspection of any good image (1.2). Their algorithm can be made faster, albeit Monte Carlo, by using randomness; we probabilistically search for a good prime by selecting primes at random over a specified range, choosing as our good prime p one for which the image $f \bmod (z^p - 1)$ has maximally many terms.

An information-theoretic lower bound on the total probe size required is $\Omega(T(\log D + \log q))$ bits, the number of bits used to encode f in (1.1). This bound is met by Prony's (1795) original algorithm, which requires a total probe size of $O(T \log q)$ under the implicit assumption that $q > D$. Much of the complexity of the sparse interpolation problem appears to arise from the requirement to accommodate *any* finite field. Prony's algorithm is dominated by the cost of discrete logarithms in $\mathbb{F}_q$, for which no polynomial time algorithm is known in general. When there is a choice of fields (say, as might naturally arise in a modular scheme for interpolating integer or rational polynomials) more efficient interpolation methods have been developed. Kaltofen (1988) demonstrates a method for sparse interpolation over $\mathbb{F}_p$ for primes p such that $p-1$ is smooth; see (Kaltofen, 2010) for further exposition. In our notation, this algorithm would require an essentially optimal $\widetilde{\mathcal{O}}(LT(\log D + \log q))$ bit operations. Parallel algorithms and implementations for this case of chosen characteristic are also given by Javadi and Monagan (2010). Moreover, our need for a bound T on the number of non-zero terms is motivated by the hope of a Las Vegas or deterministic algorithm. The early termination approach of Kaltofen and Lee (2003) identifies t with high probability at no asymptotic extra cost.

1.2.1 Diversified Interpolation

We say a polynomial $g(z)$ is *diverse* if its non-zero terms have distinct coefficients. The interpolation algorithm of Giesbrecht and Roche (2011) distinguishes between images of distinct non-zero terms by *diversifying* f, that is, choosing an α such that $f(\alpha z)$ is diverse. This Monte Carlo algorithm entails three probabilistic steps. First, it determines $t = \#f$ by searching for a probable good prime. Second, it finds an α that diversifies $f \bmod (z^p - 1)$. Finally, it looks for more good primes p, and constructs the exponents e_i of f by way of Chinese remaindering on the congruences $e_i \bmod p$. This gives $f(\alpha z)$, from which it is straightforward to recover $f(z)$.

The diversified interpolation algorithm was initially described for a "remainder black box" representation of f, a weaker model than a straight-line program. To search for an appropriate $\alpha \in \mathbb{F}_q$, the algorithm requires that the field size q is greater than $T(T-1)D$. Under the SLP model we can easily adapt the algorithm to work when $q \leq T(T-1)D$; simply choose α from a sufficiently large field extension. This slight adaptation increases the cost of a probe by a

factor of $\widetilde{\mathcal{O}}((\log D + \log T)/\log q)$.

The cost of sparse SLP interpolation using the diversified algorithm is then $\widetilde{\mathcal{O}}(\log D + \log \frac{1}{\epsilon})$ probes of degree $\widetilde{\mathcal{O}}((T^2 \log D)\lceil (\log D)/(\log q)\rceil)$, where $\epsilon \in (0,1)$ is a given bound on the probability of failure. Since each operation in $\mathbb{F}_q$ costs $\widetilde{\mathcal{O}}(\log q)$ bit operations, the cost in bit operations is

$$\widetilde{\mathcal{O}}\left(LT^2 \log D(\log D + \log \tfrac{1}{\epsilon})(\log D + \log q)\right). \tag{1.3}$$

For fixed ϵ this cost becomes $\widetilde{\mathcal{O}}(LT^2 \log^2 D(\log D + \log q))$.

The diversified and probabilistic Garg-Schost interpolation algorithms may be made Las Vegas (i.e., guaranteed error free) by way of a deterministic polynomial identity testing algorithm. The fastest-known method for certifying that polynomials (given by algebraic circuits) are equal over an arbitrary field, by Bläser, Hardt, Lipton, and Vishnoi (2009), requires $\widetilde{\mathcal{O}}(LT^2 \log^2 D)$ field operations in our notation. The gains in faster interpolation algorithms would be dominated by the cost of this certification; hence they are fast Monte Carlo algorithms whose unchecked output may be incorrect with controllably small probability.

1.2.2 Recursive interpolation

The algorithm of Arnold, Giesbrecht, and Roche (2013) is faster than diversified interpolation when T asymptotically dominates either $\log D$ or $\log q$. The chief novelty behind that algorithm is to use smaller primes p with relaxed requirements. Instead of searching for good primes separating all of the non-zero terms of f, we probabilistically search for an *ok prime* p which separates *most* of the non-zero terms of f. Given this prime p we then construct images of the form $f \bmod (z^{pq_j} - 1)$, for a set of coprime moduli $\{q_1, \ldots, q_\ell\}$ whose product exceeds D, in order to build those non-colliding non-zero terms of f.

The resulting polynomial f^* contains these terms, plus possibly a small number of *deceptive terms* not occurring in f, such that $f - f^*$ now has at most $T/2$ non-zero terms. The algorithm then updates the bound $T \leftarrow T/2$ and recursively interpolates the difference $g = f - f^*$.

The total cost of the recursive interpolation algorithm is $\widetilde{\mathcal{O}}(\log D + \log \frac{1}{\epsilon})$ probes of degree $\widetilde{\mathcal{O}}(T \log^2 D)$, for a total cost of

$$\widetilde{\mathcal{O}}\left(LT \log^2 D \left(\log D + \log \tfrac{1}{\epsilon}\right) \log q\right) \tag{1.4}$$

bit operations, which is $\widetilde{\mathcal{O}}(LT \log^3 D \log q)$ when $\epsilon \in (0,1)$ is a constant.

1.3 Outline of new algorithm and this paper

As in (Arnold et al., 2013), our new algorithm recursively builds an approximation f^*, initially zero, to the input polynomial f given by an SLP. The algorithm interpolates $g = f - f^*$ with bounds $D \geq \deg(g)$ and $T \geq \#g$. We update T as we update g.

We begin with the details of our straight-line program model in Section 2. Section 3 describes how we choose a set of $\ell \in \widetilde{\mathcal{O}}(\log D)$ ok primes $p_i \in \widetilde{\mathcal{O}}(T \log D)$, $1 \leq i \leq \ell$. Given these primes p_i, we then compute images $g_{ij} = g(\alpha_j z) \bmod (z^{p_i} - 1)$ for choices of α_j that will (probably) allow us to identify images of like terms of g. This approach relies on a more general notion of diversification. We explain how we choose the values α_j and give a probabilistic analysis in Section 4. Section 5 details how we use information from the images g_{ij} to construct at least half of the terms of g.

Section 6 describes how the algorithm iteratively builds f^*, followed by a probabilistic and cost analysis of the algorithm as a whole in Section 7. Some conclusions are drawn in Section 8.

2 Probing an SLP

In this paper we will only consider a single-input, division-free SLP $\mathcal{S}_f$ that computes a univariate polynomial f over a finite field. We can reduce multivariate $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ of total degree less than D, by way of a *Kronecker substitution*. Interpolating the univariate polynomial $f(z, z^D, \ldots, z^{D^{n-1}})$ preserves the number of non-zero terms of f and allows the easy recovery of the original multivariate terms, but increases the degree bound to D^n. See (Arnold and Roche, 2014) for recent advances on this topic.

We will *probe* $\mathcal{S}_f$, that is, execute the straight-line program on inputs of our choosing. One could naively input an indeterminant z into $\mathcal{S}_f$, and expand each resulting polynomial b_i. This sets b_L to $f(z)$. Unfortunately, intermediate polynomials b_i may have a number of non-zero terms, and be of degrees, which grow exponentially in terms of L, T, and D.

Instead, we limit the size of intermediate results by selecting a symbolic p-th root of unity $z \in \mathbb{F}_q[z]/\langle z^p - 1\rangle$ as input, for different choices of p. At each instruction we expand $b_i \in \mathbb{F}_q[z]$, and reduce modulo $(z^p - 1)$. This sets b_L to $f(z) \bmod (z^p - 1)$. Using the method of Cantor and Kaltofen (1991), the cost of one such instruction becomes $\widetilde{\mathcal{O}}(p)$ field operations. More generally we will produce images of the form $f(\alpha z) \bmod (z^p - 1)$, where α is algebraic with degree s over $\mathbb{F}_q$. That is, we choose as input αz, where z is again a p-th root of unity. Here the instruction cost becomes $\widetilde{\mathcal{O}}(p)$ operations in $\mathbb{F}_{q^s}$, or $\widetilde{\mathcal{O}}(ps)$ operations in $\mathbb{F}_q$. We call ps the *probe degree* of a probe with input αz.

We will also have to produce such images of polynomials given explicitly by sparse representations rather than implicitly in a SLP. Given a single term cz^e, $e < D$, we can produce $c(\alpha z)^e \bmod (z^p - 1)$ by computing α^e via a square-and-multiply approach and reducing the exponent of z^e modulo p. For $p < D$ the former step dominates this cost and requires $\widetilde{\mathcal{O}}(s \log D)$ operations in $\mathbb{F}_q$. For f^*, a sum of at most T such terms and $p < D$, the cost of constructing $f^*(\alpha z) \bmod (z^p - 1)$ becomes $\widetilde{\mathcal{O}}(sT \log D)$.

In particular, given an SLP $\mathcal{S}_f$ and a sparse polynomial f^*, we will need to compute $f(\alpha z) - f^*(\alpha z) \bmod (z^p - 1)$, as described above. Procedure ComputeImage

is a subroutine to perform such a computation, and its cost is $\widetilde{\mathcal{O}}(Lsp + sT \log D)$ operations in $\mathbb{F}_q$. When p is at least of magnitude $\Omega((T \log D)/L)$, the SLP probe dominates the cost, and it becomes simply $\widetilde{\mathcal{O}}(Lsp \log q)$ bit operations.

Procedure ComputeImage($\mathcal{S}_f, f^*, \alpha, p$)

Input: $\mathcal{S}_f$, an SLP computing $f \in \mathbb{F}_q[z]$; $f^* \in \mathbb{F}_q[x]$ given explicitly;
$\alpha \in \mathbb{F}_{q^s}$; integer $p \geq 1$.
Result: $f(\alpha z) - f^*(\alpha z) \bmod (z^p - 1)$

$g_{p,\alpha} \longleftarrow f(\alpha z) \bmod (z^p - 1)$, by probing $\mathcal{S}_f$ at αz over
$\mathbb{F}_{q^s}[z]/\langle z^p - 1 \rangle$;
foreach $e \in \texttt{exponents}(f^*)$ **do**
  $g_{p,\alpha} \longleftarrow g_{p,\alpha} - \texttt{coeff}(f^*, e) \cdot \alpha^e \cdot z^{e \bmod p}$
return $g_{p,\alpha}$

3 Selecting primes

The aim of Sections 3–5 is to build a polynomial f^{**} that contains at least half of the terms of $g = f - f^*$. For notational simplicity, in these sections we write $g = \sum_{i=1}^{t} c_i z^{e_i}$ and let T and D bound $\#g$ and $\deg(g)$ respectively.

We will use images of the form $g \bmod (z^p - 1)$, p a prime, in order to extract information about the terms of g. We say two terms cz^e and $c'z^{e'}$ of g *collide* modulo $(z^p - 1)$ if p divides $e - e'$, that is, if their reduced sum modulo $(z^p - 1)$ is a single term. If a term cz^e of g does not collide with any other term of g, then its image $cz^{e \bmod p}$ contained in $g \bmod (z^p - 1)$ gives us its coefficient c and the image of the exponent $e \bmod p$.

We let $\mathcal{C}_g(p)$ denote the number of terms of g that collide with any other term modulo $(z^p - 1)$. We need $\mathcal{C}_g(p)$ to be small so that the next approximation f^* of f contains many terms of f, but the primes p must also be small so that the cost of computing each f^* is not too great. In this technical section, we show how to bound the size of the primes p to balance these competing concerns and minimize the total cost of the algorithm.

In the first phase of the algorithm, we will look for a set of primes $p_i \in \widetilde{\mathcal{O}}(T \log D)$, and corresponding images

$$g_i = g \bmod (z^{p_i} - 1), \qquad 1 \leq i \leq \ell,$$

such that, in each of the images g_i, most of the terms of g are not in any collisions. More specifically, we want the primes $(p_i)_{1 \leq i \leq \ell}$ and images $(g_i)_{1 \leq i \leq \ell}$ to meet the following criteria:

(i) At least half of the terms of g do not collide modulo $(z^{p_i} - 1)$ with any other term of g for at least $\lceil \ell/2 \rceil$ of the primes p_i;

(ii) Any pair of terms of g collide modulo $(z^{p_i} - 1)$ for fewer than $\lceil \ell/2 \rceil$ of the primes p_i;

(iii) Any $\lceil \ell/2 \rceil$ of the primes p_i have a product exceeding D.

Every nonzero term of an image g_i is the image of a sum of terms of g. If we are able to collect terms of the images g_i, $1 \le i \le \ell$, that are images of the same sum of terms of g, then from (ii), any such collection containing a term for at least $\ell/2$ images g_i must be an image of a single term of g. By (i), at least half of the terms of g will produce such a collection. By (iii), such collections will contain sufficient information to reconstruct the exponent of the corresponding term of g. Thus, given a means of collecting terms in this way, we will be able to construct half of the terms of g.

We note, since any two terms of g have degree differing by at most D, that (iii) implies (ii). To satisfy (i), it suffices that $\mathcal{C}_g(p_i) \le T/2$ for each prime p_i. We will accordingly call p an *ok prime* in this paper if $\mathcal{C}_g(p) < T/2$. To this end we establish a range in which most primes p have $\mathcal{C}_g(p) < T/4$, half the desired bound.

Lemma 3.1 (Arnold et al. 2013)**.** *Let $g \in \mathbb{F}_q[z]$ be a polynomial with $t \le T$ terms and degree $d \le D$. Let*

$$\lambda = \max\left(21, \left\lceil \tfrac{40}{3}(T-1)\ln D \right\rceil\right). \tag{3.1}$$

Then fewer than half of the primes p in the range $[\lambda, 2\lambda]$ satisfy $\mathcal{C}_g(p) \ge T/4$.

We will look for ok primes in the range $[\lambda, 2\lambda]$. To satisfy (iii), we will select $\ell = 2\lceil \log_\lambda D \rceil$ primes. Corollary 3.2 gives us a means of selecting a set of primes for which a constant proportion of those primes p have fewer than $T/4$ colliding terms.

Corollary 3.2. *Suppose g, D, T, and λ are as in Lemma 3.1. Let $0 < \mu < 1$, and suppose*

$$\ell = 2\lceil \log_\lambda D \rceil, \qquad \gamma = \lceil \max(8 \log_\lambda D, 8 \ln \tfrac{1}{\mu}) \rceil. \tag{3.2}$$

If we randomly select γ distinct primes from $[\lambda, 2\lambda]$, then $\mathcal{C}_g(p) \le T/4$ for at least ℓ of the chosen primes p, with probability at least $1 - \mu$.

Proof. By Lemma 3.1, at least half of the primes p in $[\lambda, 2\lambda]$ have $\mathcal{C}_g(p) < T/4$. We require that a proportion of $\frac{1}{4}$ of the γ primes randomly selected to have this property.

In order to obtain a bound on at least $\gamma/4$ of the primes having this property, define random variables X_i, for $1 \le i \le \gamma$, to be 1 if the ith prime p_i satisfies $\mathcal{C}_g(p_i) < T/4$, and 0 otherwise. From the reasoning above, we know that $E[X_i] \ge \frac{1}{2}$. Therefore, the expected number of primes with this property, $E[\Sigma X_i]$, is at least $\frac{\gamma}{2}$.

Hoeffding's inequality provides proof that the *actual* value of ΣX_i is not too much smaller than this expected value. Specifically, Theorems 1 and 4 of (Hoeffding, 1963) show that

$$\Pr\{\Sigma X_i \le \tfrac{\gamma}{4}\} \le \exp\left(-2\left(E[\Sigma X_i] - \tfrac{\gamma}{4}\right)^2\right) \le \exp\left(\tfrac{-\gamma}{8}\right).$$

For $\gamma \ge 8 \ln \frac{1}{\mu}$, this is less than μ. ∎

We will generate some γ primes, of which ℓ primes p have $\mathcal{C}_g(p) < T/4$. In order to identify some primes p for which $\mathcal{C}_g(p)$ is low, we have the following lemma.

Lemma 3.3 (Arnold et al. (2013), Corollary 11)**.** *Suppose $g \bmod (z^q - 1)$ has s_q non-zero terms, and $g \bmod (z^p - 1)$ has s_p non-zero terms, with $s_p \geq s_q$. Then $\mathcal{C}_g(p) \leq 2\mathcal{C}_g(q)$.*

Corollary 3.2 guarantees with high probability that at least ℓ primes p_i of the γ selected satisfy $\mathcal{C}_g(p_i) < T/4$. Call these the "best primes". Unfortunately, there is no easy way to determine $\mathcal{C}_g(p_i)$ at this point, so we do not know which ℓ primes are the best primes!

Our solution to this seeming dilemma is to order the primes in decreasing order of $\#g_i$. Now consider any of the first ℓ primes p_i in this ordering. Since $\#g_i$ is among the largest ℓ values, we know that $\#g_i \geq \#g_j$ for some p_j that is one of the "best primes" and satisfies $\mathcal{C}_g(p_j) < T/4$. Then by Lemma 3.3, we know that $\mathcal{C}_g(p_i) \leq 2\mathcal{C}_g(p_j)$, which is less than $T/2$.

In other words, the first ℓ primes in this ordering are not necessarily the "best", but they are good enough because they all satisfy $\mathcal{C}_g(p_i) < T/2$. This method is described in procedure FindPrimes below.

The statement of Corollary 3.2 assumes that there are a certain number of primes in $[\lambda, 2\lambda]$, in order to pick γ of them. Letting n be the actual number of primes in this range, we require that $n \geq \gamma = \lceil \max(8 \log_\lambda D, 8 \ln \frac{1}{\mu}) \rceil$. This puts further constraints on λ. By Corollary 3 of Rosser and Schoenfeld (1962), and using the definition of λ from Lemma 3.1,

$$n \geq 3\lambda/(5 \ln \lambda) \geq 8(T-1) \log_\lambda D,$$

which is at least $8 \log_\lambda D$ for $T \geq 2$. Thus, for $T > 1$, we only require that $n \geq \lceil 8 \ln \frac{1}{\mu} \rceil$. If the number of primes n is even smaller than this, one could simply compute all n primes in the interval and use them all instead of picking a random subset. Since at least half these primes have $\mathcal{C}_g(p) < T/4$, we would only require that $n \geq 2\ell$, which must be true since $n \geq 8 \log_\lambda D$.

To ensure that computing all n primes in the interval does not increase the overall cost of the algorithm, consider that in this case γ exceeds $3\lambda/(5 \ln \lambda)$. Then the upper bound of Corollary 3 of Rosser and Schoenfeld (1962) gives

$$n \leq \lceil 7\lambda/(5 \ln \lambda) \rceil \in O(\gamma).$$

Therefore whether we choose only γ primes from the interval or are forced to compute all n of them, the number of primes used is always $O(\gamma)$.

We make one further note on the cost of producing γ primes at random. In a practical implementation, one would likely choose numbers at random (perhaps in a manner that avoids multiples of small primes), and perform Monte Carlo primality testing to verify whether such a number is prime. Problems could arise if the algorithm produced pseudoprimes p_i and p_j that are not coprime. Thus one would also have to consider the failure probability of primality testing in the analysis of such an approach.

For the purposes of our analysis, we generate all the primes up to 2λ with $\widetilde{\mathcal{O}}(\lambda) = \widetilde{\mathcal{O}}(T \log D)$ bit operations using a wheel sieve (Pritchard, 1982), which does not dominate the cost in (3.3), and guarantees that all of the chosen p are actually prime.

The total cost is γ probes of degree at most 2λ, where

$$\lambda \in \widetilde{\mathcal{O}}(T \log D), \qquad \text{and} \qquad \gamma \in \widetilde{\mathcal{O}}\left(\frac{\log D}{\log T} + \log \tfrac{1}{\mu}\right),$$

for a field-operation cost of

$$\widetilde{\mathcal{O}}(L\gamma\lambda) = \widetilde{\mathcal{O}}\left(L\left(\log D + \log \tfrac{1}{\mu}\right) T \log D\right). \tag{3.3}$$

Procedure FindPrimes($\mathcal{S}_f, f^*, T, D, \mu$)

Input: $\mathcal{S}_f$, an SLP computing $f \in \mathbb{F}_q[z]$; $f^* \in \mathbb{F}_q[x]$ given explicitly; $T \geq \max(\#g, 2)$; $D \geq \deg(g)$; $0 < \mu < 1/3$; where g is defined as the unknown polynomial $f - f^*$.

Result: A list of primes p_i and images $g(z) \bmod (z^{p_i} - 1)$ such that $\prod_{i=1}^{\ell} p_i \geq D$ and, with probability exceeding $1 - \mu$, $\mathcal{C}_g(p_i) < T/2$ for $1 \leq i \leq \ell$.

$\lambda \longleftarrow \max\left(21, \left\lceil \frac{40}{3}(T-1)\ln D \right\rceil\right)$
$\gamma \longleftarrow \lceil \max(8 \log_\lambda D, 8 \ln \frac{1}{\mu}) \rceil$
$\ell = 2\lceil \log_\lambda D \rceil$
if $\gamma \leq 3\lambda/(5 \ln \lambda)$ **then**
  $\mathcal{P} \longleftarrow$ γ primes chosen at random from $[\lambda, 2\lambda]$
else
  $\mathcal{P} \longleftarrow$ all primes from $[\lambda, 2\lambda]$
foreach $p \in \mathcal{P}$ **do** $g_p \longleftarrow$ ComputeImage($\mathcal{S}_f, f^*, 1, p$)
$(p_1, p_2, \dots) \longleftarrow \mathcal{P}$ sorted in decreasing order of $\#g_{p_i}$
return $(p_1, \dots, p_\ell), (g_{p_1}, \dots, g_{p_\ell})$

---

4 Detecting deception

At this stage we have probably found primes p_i, $\mathcal{C}_g(p_i) \leq T/2$, and their corresponding images g_i, for $1 \leq i \leq \ell$. The challenge now remains to collect terms amongst the images g_i that are images of the same term of g. For our purposes we need a more general notion of diversification than that introduced by Giesbrecht and Roche (2011). To this end we will construct images

$$g_{ij} = g(\alpha_j z) \bmod (z^{p_i} - 1), \tag{4.1}$$

where $\alpha_j \neq 0$ belongs to $\mathbb{F}_q$ or a field extension $\mathbb{F}_{q^s}$.

Any term of an image g_i is an image of either a single term of g, or a sum of multiple terms of g. Our algorithm needs to identify and discard the terms in g_i corresponding to multiple terms of g, using only the single terms to reconstruct the actual terms in g. To analyse this situation, consider the bivariate polynomials

$$g(yz) \bmod (z^{p_i} - 1) = h_{i,0} + h_{i,1}z + \cdots + h_{i,p_i-1}z^{p_i-1},$$

where each $h_{i,u} \in \mathbb{F}_{q^s}[y]$ is the sum of the terms in g with degrees congruent to u modulo p_i. Each $h_{i,u}$ has between 0 and T terms and degree at most D in y. For a given α, computing $g_i(\alpha z)$ gives the univariate polynomial whose coefficient of degree u is $h_{i,u}(\alpha)$.

Consider a single term in the unknown polynomial g. If that term does not collide with any others mod p_i, then for some u, $h_{i,u}$ consists of that single term. If the same term does not collide modulo p_k, then there exists some v such that $h_{i,u} = h_{k,v}$. Obviously, for any $\alpha \in \mathbb{F}_{q^s}$, we will have $h_{i,u}(\alpha) = h_{k,v}(\alpha)$, and our algorithm can use this correlation to reconstruct the term, since its exponent equals u mod p_i and v mod p_k. Based on the previous section, most of the terms in g will not collide modulo most of the p_i's, and so there will be sufficient information here to reconstruct those terms.

The problem is that we may have $h_{i,u}(\alpha) = h_{k,v}(\alpha)$, but $h_{i,u} \neq h_{k,v}$. We call this a *deception*, since it may fool our algorithm into reconstructing a single term in g that does not actually exist. Our algorithm will evaluate with multiple choices α_j, $1 \leq j \leq m$, for α, and we "hope" that, whenever $h_{i,u} \neq h_{k,v}$, at least one of α_j's gives $h_{i,u}(\alpha_j) = h_{k,v}(\alpha_j)$. In this case we say α_j *detects* the deception. The following lemma provides this hope.

Lemma 4.1. *Let $g \in \mathbb{F}_q[z]$ be a polynomial of degree at most D and at most T nonzero terms. Let p_i be a prime such that $\mathcal{C}_g(p_i) < T/2$ and let $g_i = g \bmod (z^{p_i} - 1)$ for $1 \leq i \leq \ell$. Let*

$$m = \left\lceil \log \tfrac{1}{2\mu} + 2\log T + 2\log(1 + \ell/4) \right\rceil, \text{ and}$$
$$s = \lceil \log_q(2D + 1) \rceil.$$

Choose $\alpha_1, \ldots, \alpha_m$ at random from $\mathbb{F}_{q^s}^*$. Then, with probability at least $1 - \mu$, every deception amongst the images $g_1, \ldots, g_\ell$ is detected by at least one of the α_j.

Proof. Consider $h_{i,u}, h_{k,v}$ as above, with $h_{i,u} \neq h_{k,v}$. As $(h_{i,u} - h_{k,v})(y)$ has degree at most D, there are at most D choices of α for which $h_{i,u}(\alpha) = h_{k,v}(\alpha)$.

Thus, if $q^s > 2D$ and $\deg(h_{i,u} - h_{k,v}) \leq D$, then at most half of the $\alpha \in \mathbb{F}_{q^s}^*$ can comprise a root of a $h_{i,u} - h_{k,v}$. Thus, if we then select $\alpha_1, \ldots, \alpha_m$ at random from $\mathbb{F}_{q^s}^*$ and construct images $g(\alpha_j z) \bmod (z^{p_i} - 1)$ for $i = 1, \ldots, \ell$, $j = 1, \ldots, m$, we will fail to detect a single given deception with probability at

most

$$\left(\frac{1}{2}\right)^m = \left(\frac{1}{2}\right)^{\lceil \log(1/(2\mu))+2\log(T)+2\log(1+\ell/4)\rceil} \leq \frac{\mu}{\frac{1}{2}T^2(1+\ell/4)^2}\,. \tag{4.2}$$

As g has at most T terms, there are at most 2^T-1 possible choices for $h_{i,u} \neq 0$; however, only a small proportion of these choices may correspond to a term in an image g_i. Each g_i, $1 \leq i \leq \ell$, contains at most $\mathcal{C}_g(p_i)/2 \leq T/4$ terms that are images of a sum of at least two terms of g. The remaining nonzero terms of g_i are images of single terms of g, of which there are at most T. Thus there are fewer than $T(1+\ell/4)$ distinct sums corresponding to one of the nonzero $h_{i,u}$.

A deception occurs between a pair of such sums, thus there are fewer than $\frac{1}{2}T^2(1+\ell/4)^2$ deceptions. It then follows from (4.2) that the probability that at least one deception is not detected by any of $\alpha_1, \ldots, \alpha_m$ is bounded above by μ. □

It is important to note that m is logarithmic in T and ℓ, so that a multiplicative factor of m will not affect the "soft-Oh" cost of the algorithm in terms of T or D. If $q \leq 2D$, then we need instead to work in an extension of $\mathbb{F}_q$ of degree $s = \lceil \log_q(2D+1) \rceil$. The cost of computing g_{ij}, $1 \leq i \leq \ell, 1 \leq j \leq m$, in terms of $\mathbb{F}_q$-operations, becomes

$$\widetilde{\mathcal{O}}(L\ell m\lambda s) = \widetilde{\mathcal{O}}\left(LT\log^2 D\left\lceil\frac{\log D}{\log q}\right\rceil \log\tfrac{1}{\mu}\right),$$

which dominates the cost (3.3) of Section 3.

5 Identifying images of like terms

As we construct the images g_{ij}, we will build vectors of coefficients of images g_{ij}. Namely, for every congruence class $e \bmod p_i$ for which there exists a nonzero term of degree e in at least one image g_{ij}, we will construct a vector $v^{i,e} \in \mathbb{F}_{q^s}^m$, where $v_j^{i,e}$ contains the coefficient (possibly zero) of the term of g_{ij} of degree e.

We will use the vectors $v^{i,e}$ to identify terms of the images g_i that are images of like terms of g. We use these vectors $v^{i,e}$ as keys in a dictionary `tuples`. Each value in the dictionary is comprised of a list of those tuples (e,i) for which $v^{i,e} = v$.

Provided the probabilistic steps of Sections 3 and 4 succeeded, if a key v is found more than $\ell/2$ times, then it corresponds to a single, distinct term of g, as opposed to a sum of terms of g. This is indicated by the size of the list `tuples`(v) being at least $\ell/2$.

The dictionary `tuples` should be an ordered dictionary that supports logarithmic-time insertion and retrieval. Any balanced binary search tree, such as a red-black tree will be suitable (see, e.g., (Cormen, Leiserson, Rivest, and Stein, 2001), Chapter 13). To set `tuples`(v), we first search to see if v is an existing key; if not, an empty list is first inserted as the value of `tuples`(v).

A red-black tree of size n requires $\mathcal{O}(\log n)$ comparisons for insert and search operations. We compare keys $v \in \mathbb{F}_{q^s}^m$ lexicographically, which may entail m comparisons of elements in $\mathbb{F}_{q^s}$. Each comparison therefore requires $\mathcal{O}(ms \log q)$ bit operations.

The number of different vectors $v^{i,d}$ that will appear is bounded above by the number of distinct subsets of terms of g which can collide modulo $(z^{p_i} - 1)$, for any of the primes p_i. Since there are ℓ primes, and at most T terms in g, there are no more than $T\ell$ vectors which will be inserted as keys into `tuples`. Thus, the cost of constructing this tree is

$$\begin{aligned} &\mathcal{O}(T\ell ms \cdot \log q \cdot \log(T\ell)) \\ &\quad = \widetilde{\mathcal{O}}\left(T \log D \, (\log D + \log q) \log \tfrac{1}{\mu}\right) \end{aligned} \tag{5.1}$$

bit operations. This cost is dominated by that of constructing the g_{ij}. Each term in each g_{ij} contains an element of $\mathbb{F}_{q^s}$ and an exponent at most 2λ. This requires $\widetilde{\mathcal{O}}(s \log q + \log \gamma)$ bits, which is $\widetilde{\mathcal{O}}(\log D + \log q + \log T + \log \frac{1}{\mu})$. The additional bit-cost of traversing the images g_{ij} and appending to the lists in $\texttt{tuples}(v)$ is reflected in the cost of their construction in Section 4.

After we have constructed the dictionary `tuples`, we traverse it again to build terms of g. For every key v whose corresponding list has size at least $\ell/2$, we have all the pairs (i, d) such that $v^{i,d} = v$. What remains is to construct the term corresponding to the key v. We reconstruct the exponent by Chinese remaindering on the first $\ell/2$ congruences $d \bmod p_i$. As each exponent is at most D, the cost of constructing one exponent is bounded by $\widetilde{\mathcal{O}}(\log^2 D)$ bit operations. Thus the total bit-cost of Chinese remaindering becomes $\widetilde{\mathcal{O}}(T \log^2 D)$. As the bit cost of s operations in $\mathbb{F}_q$ is $\widetilde{\mathcal{O}}(\log q + \log D)$, this cost of Chinese remaindering is bounded asymptotically by (5.1), the cost of constructing `tuples` itself.

We obtain the coefficient by inspection of g_i. We sum all of these constructed terms into a polynomial f^{**} approximating $g = f - f^*$.

Procedure BuildApproximation restates the method described to construct f^{**}. For the sake of brevity, this procedure does not detail the data structures used for the polynomials g_{ij} and f^{**} that it constructs. In order to achieve the stated complexity bounds, these sparse polynomials must be implemented by dictionaries mapping exponents to coefficients that support logarithmic-time insertion and retrieval. As with `tuples`, a red-black tree, hash table, or similar standard data structure could be used. Converting between such a representation and the usual list of coefficient-exponent pairs is trivial.

Procedure BuildApproximation($\mathcal{S}_f, f^*, T, D, \mu$)

Input: $\mathcal{S}_f$, an SLP computing $f \in \mathbb{F}_q[z]$; $f^* \in \mathbb{F}_q[x]$ given explicitly; $T \geq \#g$; $D \geq \deg(g)$; $0 < \mu < 1/3$; where g is defined as the unknown polynomial $f - f^*$.

Result: f^{**} such that $g - f^{**}$ has at most $T/2$ terms, with probability greater than $1 - \mu$.

$(p_1, \ldots, p_\ell), (g_1, \ldots, g_\ell) \longleftarrow$ FindPrimes($\mathcal{S}_f, f^*, T, D, \mu$)
$m \longleftarrow \left\lceil \log \frac{1}{2\mu} + 2 \log T + 2 \log(1 + \ell/4) \right\rceil$
$s \longleftarrow \lceil \log_q(2D + 1) \rceil$
$(\alpha_1, \ldots, \alpha_m) \longleftarrow m$ randomly chosen nonzero elements from $\mathbb{F}_{q^s}$;
`tuples` $\longleftarrow$ dictionary mapping $\mathbb{F}_{q^s}^m$ to lists of pairs of integers;

for $i \longleftarrow 1$ **to** ℓ **do**
  foreach $e \in \bigcup_{1 \leq j \leq m}$ `exponents`(g_{ij}) **do**
    $v \longleftarrow \big($`coeff`$(g_{i1}, e), \ldots,$ `coeff`$(g_{im}, e)\big)$
    `tuples`$(v) \longleftarrow$ `tuples`(v), (i, e)

$f^{**} \longleftarrow 0$
foreach $v \in$ `tuples` **do**
  if #`tuples`$(v) \geq \ell/2$ **then**
    $e \longleftarrow$ solution to congruences $\{e_i \bmod p_i \mid (i, e) \in$ `tuples`$(v)\}$;
    $c \longleftarrow$ `coeff`(g_i, e), for any of the $(i, e) \in$ `tuples`(v);
    $f^{**} \longleftarrow f^{**} + cz^e$

return f^{**}

6 Updating our approximation

Recall we have a polynomial f^* approximating f given by our straight-line program. We construct a polynomial f^{**} that is comprised of at least $T/2$ terms of $g = f - f^*$. Once we have f^{**}, we set $f^* \leftarrow f^* + f^{**}$, $T \leftarrow \lfloor T/2 \rfloor$, and repeat the process until T is 1, at which point g consists of (at most) a single nonzero term.

We thus execute this process at most $\log T$ times, where T is the initial bound on the number of non-zero terms of f. Recall that the steps of sections 3 and 5 each succeed with probability greater than $1 - \mu$. Thus, if we would like the algorithm to succeed with probability greater than $1 - \epsilon > 1/2$, we can set

$$\mu = \epsilon/(2 \log T). \tag{6.1}$$

As $T \geq 2$ and $\epsilon \leq 1/2$, we will always have $\mu \leq 1/4$, which satisfies the constraint

$\mu < 1/3$ from Section 3.

When $T = 1$, g is a single term and its coefficient is given by $g(1)$. The exponent of the term comprising g may be computed from $g \bmod (z^p - 1)$ for the first $\log D$ primes p. This cost is $\log D$ probes of degree $\widetilde{\mathcal{O}}(\log D)$, or a cost of $\widetilde{\mathcal{O}}(\log^2 D \log q)$ bit operations.

We give the interpolation algorithm in procedure MajorityVoteSparseInterpolate. We call the algorithm "majority-vote" sparse interpolation, as we effectively require a majority of the images f_i, $1 \le i \le \ell$, to vote on whether a sum of terms of f is in fact a single term.

Procedure MajorityVoteSparseInterpolate($\mathcal{S}_f, T, D, \epsilon$)

Input: $\mathcal{S}_f$, an SLP computing $f \in \mathbb{F}_q[z]$; $T \ge \#f$; $D \ge \deg(f)$;
$0 < \epsilon < 1/2$.

Result: The sparse representation of f, with probability at least $1 - \epsilon$.

$(f^*, \mu) \longleftarrow (0,\ \epsilon/(2 \log T))$
while $T > 1$ **do**
  $f^{**} \longleftarrow$ BuildApproximation($\mathcal{S}_f, f^*, T, D, \mu/2$)
  $f^* \longleftarrow f^* + f^{**}$
  $T \longleftarrow \lfloor T/2 \rfloor$

$(c, e) \longleftarrow$ (ComputeImage($\mathcal{S}_f, f^*, 1, 1$), 0)
if $c = 0$ **then return** f^* **for** *the first* $\log D$ *primes* p **do**
  $e' \longleftarrow$ degree of the single term in
    ComputeImage($\mathcal{S}_f, f^*, 1, p$);
  Update e with e' modulo p by Chinese remaindering.
return $f^* + cz^e$

7 Cost analysis

We now are ready to analyze the cost of the algorithm MajorityVoteSparseInterpolate and verify Theorem 1.1. As we have argued in Sections 3–5, the cost of one iteration of the algorithm is dominated by the cost of constructing the images g_{ij}. Recall that there are ℓ primes p_i, each less than or equal to 2λ, and there are m elements α_j's, each in an extension $\mathbb{F}_{q^s}$. The total number of field operations is thus $\widetilde{\mathcal{O}}(L\ell m\lambda s)$. The values of L, T, D, and ϵ are specified in the input. Recall the following parameters from equation (6.1), Corollary 3.2 and Lemma 4.1:

$$
\begin{aligned}
\mu &= \epsilon/(2\log T) \Rightarrow \log \tfrac{1}{\mu} \in \widetilde{\mathcal{O}}(\log \tfrac{1}{\epsilon} + \operatorname{loglog} T), \\
\lambda &= \max\left(21, \left\lceil \tfrac{40}{3}(T-1)\ln D \right\rceil\right) \in \widetilde{\mathcal{O}}(T \log D), \\
\ell &= 2\lceil \log_\lambda D \rceil \in \widetilde{\mathcal{O}}((\log D)/(\log T)), \\
m &= \left\lceil \log \tfrac{1}{2\mu} + 2\log T + 2\log(1+\ell/4) \right\rceil, \\
&\in \widetilde{\mathcal{O}}(\log \tfrac{1}{\epsilon} + \log T + \operatorname{loglog} D), \\
s &= \lceil \log_q(2D+1) \rceil \in \widetilde{\mathcal{O}}(1 + (\log D)/(\log q)).
\end{aligned}
$$

Therefore the total cost of the $\widetilde{\mathcal{O}}(L\ell m\lambda s)$ operations in $\mathbb{F}_q$ is

$$
\widetilde{\mathcal{O}}\left(LT \log^2 D \,(\log D + \log q) \log \tfrac{1}{\epsilon}\right) \tag{7.1}
$$

bit operations. The multiplicative $\log T$ factor due to the number of iterations does not affect the "soft-Oh" cost above. This cost is a multiplicative factor of $\widetilde{\mathcal{O}}(T)$ improvement over the cost (1.3) of diversified interpolation (Giesbrecht and Roche, 2011). It also improves over the cost (1.4) of recursive interpolation (Arnold et al., 2013) by a multiplicative factor of $\widetilde{\mathcal{O}}(\min(\log D, \log q))$.

8 Conclusions

We have presented a new algorithm for sparse interpolation over an arbitrary finite field that is asymptotically faster than those previously known. In terms of bit operations it improves on previous methods by "soft-Oh" multiplicative factor of T, $\log D$, or $\log q$.

In this "majority-vote" interpolation, we combine the main ideas of the previous two algorithms. Namely, we reduce the probe degree by only aiming to reconstruct *some* of the terms of f at every iteration; and we distinguish images of distinct terms (and subsets of terms) by way of diversification.

We mention a few open and motivating problems. First, we believe the algorithm in this paper has considerable potential in the more traditional numerical (floating point) domain. There, an unknown sparse polynomial is reconstructed from a small number of evaluations in $\mathbb{C}$ under a standard backward-error model of precision and stability. See (Giesbrecht and Roche, 2011) for an example of a straight-line program interpolation algorithm adapted to floating point computation. Our hope is that by evaluating at lower-order roots of unity (such as used in this paper) we can provably increase the numerical stability over the Prony-like algorithm of Giesbrecht, Labahn, and Lee (2009), while maintaining its near-optimal efficiency.

Second, the algorithm presented here is Monte Carlo. It remains unknown whether there exist faster deterministic or Las Vegas polynomial identity tests that may render recursive or majority-vote interpolation Las Vegas of the same complexity.

Finally, as noted earlier, an information theoretic lower-bound on sparse interpolation suggests a minimum bit-complexity of $\widetilde{\mathcal{O}}(LT(\log q + \log D))$ bit operations. While this paper gets closer, some considerable improvements remain to be found.

9 Acknowledgements

The first two authors would like to acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). The third author is supported by the National Science Foundation, award #1319994.

References

Andrew Arnold and Daniel S. Roche. Multivariate sparse interpolation using randomized Kronecker substitutions. To appear, International Symposium on Symbolic and Algebraic Computation (ISSAC'14), 2014.

Andrew Arnold, Mark Giesbrecht, and Daniel S. Roche. Faster sparse interpolation of straight-line programs. In *Proceedings of Computer Algebra in Scientific Computation (CASC 2013)*, volume 8136 of *Lecture Notes in Computer Sciences*, pages 61–75, 2013.

Markus Bläser, Moritz Hardt, Richard J. Lipton, and Nisheeth K. Vishnoi. Deterministically testing sparse polynomial identities of unbounded degree. *Information Processing Letters*, 109(3):187–192, 2009.

David G. Cantor and Erich Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28:693–701, 1991.

Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. The MIT Press, second edition, September 2001.

Timothy S. Freeman, Gregory M. Imirzian, Erich Kaltofen, and Lakshman Yagati. Dagwood: A system for manipulating polynomials given by straight-line programs. *ACM Trans. Math. Softw.*, 14:218–240, September 1988. doi: 10.1145/44128.214376.

Sanchit Garg and Éric Schost. Interpolation of polynomials given by straight-line programs. *Theor. Comput. Sci.*, 410(27-29):2659–2662, June 2009. ISSN 0304-3975.

Mark Giesbrecht and Daniel S. Roche. Diversification improves interpolation. In *Proceedings of the 36th international symposium on Symbolic and algebraic computation*, ISSAC '11, pages 123–130, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0675-1.

Mark Giesbrecht, George Labahn, and Wen-shin Lee. Symbolic-numeric sparse interpolation of multivariate polynomials. *Journal of Symbolic Computation*, 44:943–959, 2009.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):pp. 13–30, 1963. ISSN 01621459. URL http://www.jstor.org/stable/2282952.

Seyed Mohammad Mahdi Javadi and Michael Monagan. Parallel sparse polynomial interpolation over finite fields. In *Proc. International Workshop on Parallel Symbolic Computation (PASCO 2010)*, pages 160–168, 2010.

Erich Kaltofen and Wen-shin Lee. Early termination in sparse interpolation algorithms. *J. Symbolic Comput.*, 36(3-4):365–400, 2003.

Erich L. Kaltofen. Unpublished article fragment. http://www.math.ncsu.edu/kaltofen/bibliography/88/Ka88_ratint.pdf, 1988.

Erich L. Kaltofen. Fifteen years after DSC and WLSS2 What parallel computations I do today. In *Proc. International Workshop on Parallel Symbolic Computation (PASCO 2010)*, pages 10–17, 2010.

Paul Pritchard. Explaining the wheel sieve. *Acta Informatica*, 17(4):477–485, 1982.

Baron de Prony, Gaspard-Clair-François-Marie Riche. Essai expérimental et analytique sur les lois de la Dilatabilité des fluides élastique et sur celles de la Force expansive de la vapeur de l'eau et de la vapeur de l'alkool, à différentes températures. *J. de l'École Polytechnique*, 1:24–76, 1795.

J. Barkley Rosser and Lowell Schoenfeld. Approximate formulas for some functions of prime numbers. *Illinois J. Math.*, 6:64–94, 1962. ISSN 0019-2082.

Volker Strassen. Algebraic complexity theory. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, pages 633–672. MIT Press, 1990.